

Arrhythmia Classification from the Abductive Interpretation of Short Single-Lead ECG Records

Tomás Teijeiro*, Constantino A. García, Daniel Castro and Paulo Félix

Centro Singular de Investigación en Tecnoloxías da Información (CITIUS), University of Santiago de Compostela, Santiago de Compostela, Spain

Abstract

In this work we propose a new method for the rhythm classification of short single-lead ECG records, using a set of high-level and clinically meaningful features provided by the abductive interpretation of the records. These features include morphological and rhythm-related features that are used to build two classifiers: one that evaluates the record globally, using aggregated values for each feature; and another one that evaluates the record as a sequence, using a Recurrent Neural Network fed with the individual features for each detected heartbeat. The two classifiers are finally combined using the stacking technique, providing an answer by means of four target classes: Normal sinus rhythm (N), Atrial fibrillation (A), Other anomaly (O) and Noisy (~). The approach has been validated against the 2017 Physionet/CinC Challenge dataset, obtaining a final score of 0.83 and ranking first in the competition.

1. Introduction

The potential of Artificial Intelligence and machine learning techniques to improve the early detection of cardiac diseases using low-cost ECG tests is still largely untapped. The 2017 Physionet/Computing in Cardiology challenge defies the scientific community to propose solutions to the automatic detection of Atrial Fibrillation from short single lead ECG signals [1]. The challenge is posed as a classical machine learning problem: A labeled training set is provided, and the proposals are evaluated against a hidden test set of records. However, even if the only metric for the final ranking is the accuracy of the proposed models, a number of additional properties should be considered for the final adoption of each proposal in the clinical practice. Here, we emphasize on the interpretability of the automatic detection of Atrial Fibrillation, a major concern to ensure trust by the care staff [2].

In this sense, our proposal is based on a high-level description of the target signal by means of the same features used by cardiologists in ECG analysis. This description is generated with a pure knowledge-based approach, using

an abductive framework for time series interpretation [3] that looks for the set of explanatory hypotheses that best account for the observed evidence. Only after this description has been built, machine learning methods were used to make up for the lack of the expert criteria applied in the labeling of the training set, and to alleviate the effect of possible errors in the interpretation process.

2. Methods

The global architecture of the proposal is depicted in Figure 1, and the processing stages are explained in the following subsections.

2.1. Preprocessing

The preprocessing stage aims at improving the quality of the data to be interpreted in the following stages, and involves two different tasks:

2.1.1. Data relabeling: The labeling of the training set was performed by a single expert in a single pass, and as a consequence some inconsistencies appear in the classification criteria. Thus, a thorough manual relabeling was carried out, but trying to be conservative and guided by preliminary classification results. We focused on records classified as N but showing what we consider clear anomalies. A total number of 197 out of 8528 records were relabeled.

2.1.2. Lead inversion detection: A number of records in the training set were found to be inverted, probably due to electrode misplacement. Inverted records are more likely to be classified as abnormal due to the presence of infrequent QRS and T wave morphologies, as well as to the greater difficulty to identify P waves. The inverted records were first identified manually, and then a simple logistic regression classifier was trained considering 14 features obtained from the raw signal and a tentative delineation of the P wave, QRS complex and T wave of every heartbeat detected by the `gqrs` application from the Physionet library [4]. This delineation was performed using the *Construe* algorithm [3], limiting the interpretation to the conduction level, that is, avoiding the rhythm interpretation.




Figure 1. Classification algorithm steps.

2.2. Abductive interpretation

The abductive interpretation of the ECG signal is the most significant stage in the proposed approach. Its objective is to characterize the physiological processes underlying the signal behavior, building a description of the observed phenomena in multiple abstraction levels. This responsibility lies with the *Construe* algorithm, which applies a non-monotonic reasoning scheme to find the set of hypotheses that best explain the observed evidence, by means of a domain-specific knowledge base composed of a set of *observables* and a set of *abstraction grammars*. The knowledge base is the same used in [3], that allows to explain the ECG at the conduction and rhythm abstraction levels, thus providing the same features used by cardiologists in ECG analysis. The initial evidence is the set of waves identified in the wave delineation step, that are abstracted by a set of rhythm patterns to describe the full signal as a sequence of cardiac rhythms, including normal sinus rhythms, bradycardias, tachycardias, atrial fibrillation episodes, etc. The non-monotonic nature of the interpretation process allows us to modify the initial set of evidence, by discarding heartbeats that cannot be abstracted by any rhythm pattern, or by looking for missed beats that are predicted by the pattern selected as the best explanatory hypothesis for a signal fragment. This ability to correct the initial evidence is the main strength of our proposal, since it discards many false anomalies generated by the presence of noise and artifacts in the signal. Figure 2 shows an example of a noisy signal in which the *gqrs* application detects many false positive beats, that are removed or modified in the final interpretation that concludes with a single normal rhythm hypothesis that explains the full fragment. As we can also see in the Figure, the result of the interpretation stage is a sequence of P waves, QRS complexes and T waves observations, as well as a sequence of cardiac rhythms abstracting all those waves.

2.3. Global feature extraction

Considering that each ECG record has to be classified globally, providing a single label for the entire signal duration, after the interpretation stage a set of features are calculated trying to summarize the information provided by *Construe*. A total number of 79 features are calculated, that are comprehensively described in the published software documentation. The feature set is divided into three

main groups:

- *Rhythm features*: This includes statistical measures on the RR sequence, such as the limits, median or median absolute deviation; heart rate variability features such as the PNN5, PNN10, PNN50 and PNN100 measures [5]; and information about the rhythm interpretation, such as the median duration of each rhythm hypothesis.
- *Morphological features*: This includes information about the duration, amplitude and frequency spectrum of the observations in the conduction abstraction level, including P and T waves, QRS complexes, PR and QT intervals, and the TP segments.
- *Signal quality features*: Their purpose is to assess the importance of the morphological features showing conduction anomalies, such as wide QRS complexes or long PR intervals. They are based on the sum of the absolute differences of the signal, which we refer to as *profile*. Some of the profiled areas of the signal are the baseline segments and the P wave area before each heartbeat (taking a constant window of 250 ms).

2.4. Global classification

If a precise definition of the expert knowledge leading to the labeling of the training set were available, then the final classification could be directly developed with a basic rule-based system operating on the features extracted from the abductive interpretation stage, and the accuracy of the system would depend mainly on the accuracy of the interpretations. However, the challenge does not publish any guidelines for the classification, specially for the **O** class. Therefore, an automatic classifier was trained with two objectives: 1) To reveal the criteria leading to the training set labeling; and 2) to make the classification more accurate by learning possible mistakes of the abductive interpretation.

The classification method selected for this stage was the Tree Gradient Boosting algorithm, and particularly the XGBoost implementation [6], which showed a high performance and a certain level of interpretability through the importance given to the classification features. The optimization of the hyperparameters was performed using exhaustive grid search and 8-fold cross-validation, leading to the following values: Maximum tree depth: 6, Learning rate: 0.2, Gamma: 1.0, Column subsample by tree: 0.9, Min. child weight: 20, Subsample: 0.8, and Number of boosting rounds: 60.




Figure 2. How the abductive interpretation can fix errors in the initial evidence. [Source: First 10 seconds of the A02080 record. Grey: Original `gqrs` annotations. Blue: QRS observations. Yellow: T wave observations. Green: P wave observation and Normal rhythm hypothesis.]

With respect to the first objective, we were able to formalize a number of specific anomalies that lead to classify a record as **O**. This identification helped to optimize the training set by defining more specific features to be calculated from the interpretation results. Some of the identified anomalies sharing this class were:

- Tachycardia (Mean heart rate over 100 bpm).
- Bradycardia (Mean heart rate under 50 bpm).
- Wide QRS complex (Longer than 110 milliseconds).
- Presence of ventricular or fusion beats.
- Presence of at least one extrasystole.
- Long PR interval (Longer than 210 milliseconds).
- Ventricular tachycardia.
- Atrial flutter.

For some of these anomalies the classification in the training set seems a bit inconsistent, since examples can be found in several classes. For example, there are various records labeled as normal with PR interval longer than 210 milliseconds, as long as examples of records labeled as atrial fibrillation showing an atrial flutter pattern.

Regarding the second objective, even after discovering some of the expert criteria distinguishing the target classes a rule-based system was not still competitive against automatic learned models. From our point of view this shows that the XGBoost classifier is able to improve the results of the interpretation alone.

2.5. Per-beat feature extraction

Some of the conditions leading to a certain classification may not be present for the entire duration of a record, so the global features are not the best option to characterize episodic events of abnormalities. For example, a normal record with a single ectopic ventricular beat that does not break the rhythm is quite difficult to classify as abnormal by the global classifier. For this reason, some of the features calculated from the abductive interpretation are disaggregated to the individual heartbeat scope, such as the morphology, duration and amplitude of the P wave, the QRS complex and the T wave. Also the RR interval and the RR variation before and after each beat is included, as long as the profile of the P wave area. A sequence classification approach is then used to learn characteristic temporal patterns of each target class.

2.6. Sequence classification

In the proposed approach, sequence classification relies on Recurrent Neural Networks (RNNs), a family of neural networks specialized for recognizing sequences of values. Among the different RNN implementations, we focused on Long Short Term Memory networks (LSTMs) [7], since they are capable of remembering information for long periods of time through the use of a cell state. Furthermore, they are able to avoid vanishing and exploding gradients when doing backpropagation through time. The architecture of the neural net is shown in Figure 3. The time-distributed Multilayer Perceptron (MLP) preprocesses the features described in Section 2.5 to transform the data into a space with easier temporal dynamics. The number of hidden units of the MLP was 256, and the dimension of the output space 128. A Rectified Linear Unit (ReLU) was used as activation function. The *LSTM_0* layer preprocesses the resulting sequence of transformed features and returns a new sequence, which is subsequently used by the other LSTMs. The *LSTM_2* layer just returns the final state of the network, whereas *LSTM_1* and *LSTM_3* return new transformed sequences. The pooling layers after *LSTM_1* and *LSTM_3* remove the temporal dimension by computing the temporal mean and maximum of each feature of the sequences, respectively. All the LSTMs used 128 units. Another MLP (with the same configuration of the time-distributed one) joins and transforms the outputs of each LSTM before a Softmax layer, which outputs a probability for each of the 4 classes. L^2 -regularization was applied to all layers, using 10^{-4} as regularization strength. Finally, dropout was also used to improve generalization by preventing feature co-adaptation [8].

The neural network was trained using the categorical cross-entropy as loss function, a batch size of 32, and Adam [9] as optimizer. Furthermore, 15% of all the data was used as validation set to monitor the performance of the neural network. This permitted us to decrease the learning rate when the validation loss got stuck in a plateau and to avoid overfitting by using early stopping. The initial learning rate was set to 0.002 and it was decreased by $\sqrt{2}$ when the validation loss did not improve for at least 3 epochs. Training was ended after 15 epochs without improvement.




Figure 3. The neural network architecture.

2.7. Classification stacking

The XGBoost classifier based on global features and the RNN classifier based on the per-beat features were combined using the stacking technique. Stacking (also referred to as stacked generalization) involves training a new classification algorithm to combine the predictions of several classifiers [10]. Usually, the stacked model achieves better performance than the individual models due to its ability to discern when each base model performs best and when it performs poorly. Prior to the application of stacking, the predictions of 3 RNNs were averaged to decrease the variance of the RNN classifier arising from the random initialization of the RNN weights and the random split between test and validation set. Averaging similar models also helps in reducing overfitting. Note that this averaging can be seen as a simple bagging method. The probabilities predicted by the XGBoost and the averaged RNNs are then combined through a Linear Discriminant Analysis (LDA) classifier, which acts as stacker. To avoid possible collinearity issues, only 3 probabilities from each model are used.

3. Evaluation

To evaluate the performance of the algorithm, we followed the challenge guidelines and metrics. The final score is assigned as the mean F1 measure of the **N**, **A**, and **O** classes. Table 1 shows an example of the results that the proposed method is able to achieve using 8-fold cross-validation. Note that the stacker usually achieves better scores than the base models and, furthermore, it has lower variance (not shown in the Table).

Table 1. Example of stratified 8-fold cross-validation.

Method	Fold Number								Mean
	0	1	2	3	4	5	6	7	
XGBoost	0.84	0.84	0.85	0.85	0.82	0.80	0.82	0.82	0.83
RNN	0.82	0.81	0.84	0.83	0.86	0.83	0.83	0.83	0.83
LDA-stacker	0.85	0.84	0.86	0.86	0.85	0.83	0.84	0.85	0.85

4. Conclusions

This work proves that the combination of knowledge-based and learning-based approaches is effective to build classification systems that exploit sophisticated machine learning methods while maintaining a remarkable degree

of interpretability by the use of high-level and meaningful features.

Acknowledgements

This work was supported by the Spanish Ministry of Economy and Competitiveness under project TIN2014-55183-R. Constantino A. García is also supported by the FPU Grant program from the Spanish Ministry of Education (MEC) (Ref. FPU14/02489).

References

- [1] Clifford G, Liu C, Moody B, Lehman L, Silva I, Li Q, Johnson A, Mark R. AF classification from a short single lead ECG recording: The Physionet Computing in Cardiology Challenge 2017. In *Computing in Cardiology*. 2017; .
- [2] Caruana R, Lou Y, Gehrke J, Koch P, Sturm M, Elhadad N. *Intelligible Models for HealthCare*. In *21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM Press, 2015; .
- [3] Teijeiro T, Félix P, Presedo J, Castro D. Heartbeat classification using abstract features from the abductive interpretation of the ECG. *IEEE Journal of Biomedical and Health Informatics* 2016;.
- [4] Goldberger A, et al. *PhysioBank, PhysioToolkit and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals*. *Circulation* 2000;101:215–220.
- [5] Mietus JE, Peng CK, Henry I, Goldsmith RL, Goldberger AL. The pNNx files: re-examining a widely used heart rate variability measure. *Heart British Cardiac Society* oct 2002; 88(4):378–80. ISSN 1468-201X.
- [6] Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. In *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. mar 2016; .
- [7] Hochreiter S, Schmidhuber J. Long short-term memory. *Neural computation* 1997;9(8):1735–1780.
- [8] Srivastava N, Hinton GE, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research* 2014;15(1):1929–1958.
- [9] Kingma D, Ba J. Adam: A method for stochastic optimization. *arXiv preprint arXiv14126980* 2014;.
- [10] Wolpert DH. Stacked generalization. *Neural networks* 1992;5(2):241–259.

Address for correspondence:

Tomás Teijeiro Campo
 Rúa de Jenaro de la Fuente Domínguez, S/N, CITIUS Building
 15782 Santiago de Compostela, SPAIN
 tomas.teijeiro@usc.es